

AMENDMENTS TO THE SPECIFICATION:

Please add the following new paragraphs after paragraph beginning at page 2, line 12, as follows:

United States Patent number US6625161, entitled “ADAPTIVE INVERSE MULTIPLEXING SYSTEM AND METHOD” by Su et al, describes a method and system of combining a plurality of parallel communications channels to emulate a single high-bandwidth communication channel. A continuous stream of packets are grouped as traffic aggregates and assigned to queues associated with the plurality of parallel communication channels. The assignment and reassignment of traffic aggregates to the queues is performed dynamically based on measuring the queue load ratios associated with the lengths of the queues for each of the parallel communication channels. Grouping of existing and future packets as traffic aggregates is based on common attributes shared by the packets such as common source and destination IP addresses. The term “aggregation” as used by Su et al means that packets are aggregated within a single parallel communication channel.

United States Patent Application number US2003/021266, entitled “SCHEDULING THE DISPATCH OF CELLS IN NON-EMPTY VIRTUAL OUTPUT QUEUES OF MULTISTAGE SWITCHES USING A PIPELINED HIERARCHICAL ARBITRATION SCHEME” by Oki et al, describes how pipe-lined based matching scheduling for input-buffered switches relaxes the timing constraint for arbitration by providing a scheduling scheme for multi-stage switch arrangement. Oki et al describe a $kn \times kn$ switch arrangement having k input modules, each of the input modules comprising a group of n input ports and k output modules, each output module

comprising a group of n output ports. A number m of central modules are arranged between the input modules and the output modules, such that a three-stage switch is provided. Oki et al describe how each input module has m outgoing links L_i , each link L_i connecting the input module to a different one of the m central modules, and each of the m central modules has k outgoing links L_o , each link L_o connecting the central module to a different one of the k output modules.

Within each of the k input modules, a number of virtual output queues (VOQs) are provided, each associated with one of the nk output ports. Each input module has n input ports, resulting in each of the nk VOQs being capable of receiving in any given timeslot at most n cells from n input ports. Each VOQ is able to send just one cell to a central module in one cell timeslot. Within each of the input modules the nk VOQs are grouped into k groups of n VOQs.

Oki et al describe a multiple phase cell dispatch scheme in which VOQs of an input module and outgoing links of the input module are matched in a first phase, and then an outgoing link of an input module is matched to an outgoing link of a central module in a second phase. A hierarchical scheme is also described in which for each outgoing link of an input module, a master arbiter selects a group of VOQs from among a number of candidate groups, and a slave arbiter selects a VOQ from amongst the n VOQs belonging to the selected group.

Please replace the paragraphs beginning at page 2, line 19 through page 3, line 14, as follows:

A first aspect of the invention provides a channel assignment process for a switch arrangement, the switch arrangement being arranged to perform the switching of traffic in both the space and time domains, the process comprising:

receiving traffic along a number of ingress subelements of the switch;

for each ingress aggregation comprising a subset of said plurality of ingress subelements, aggregating one or more time-slots of the ingress subelements comprising said ingress aggregation to form one or more aggregated time-space channels;

assigning one or more inner time-space channels available through at least one inner time-shared spatial switching stage of the switch arrangement to each aggregated time-space channel; and

assigning a plurality of end-to-end time-space channels from the plurality of ingress subelements to a plurality of egress subelements through the switch arrangement after said aggregated channels have been assigned through said at least one inner time-shared spatial switching stage.

In one embodiment, each ingress aggregation further comprises a one or more switching stages forming a switching sub-structure of the switch arrangement.

In one embodiment, said switch arrangement further comprises a plurality of egress aggregations, each egress aggregation comprising a subset of the egress subelements of the switch arrangement and one or more switching stages forming a switching substructure of the switch arrangement, wherein for each egress aggregation, one or more time-slots of the egress subelements are aggregated to form one or more aggregated time-space channels.

In one embodiment, at least one switching stage comprises a time-domain switching stage.

In one embodiment, in said step of assigning a plurality of end-to-end channels, traffic received at an ingress subelement is assigned to an end-to-end channel comprising: at least one channel through an ingress aggregation; at least one channel through each of the at least one inner time-shared spatial switching stages of the switch arrangement; and at least one channel through an egress aggregation.

In one embodiment, the number of time-space inner channels provided by the logical switches of each of said at least one inner time-shared spatial switching stage of the switch arrangement to each ingress aggregation is equal to at least the number of time-slots that each ingress subelement switches end to end through the switch arrangement multiplied by the number of ingress subelements which form an ingress aggregation.

In one embodiment, the channel assignment process is implemented as a frame-based channel assignment process.

In one embodiment, said process is further iteratively performed within each ingress aggregation to determine a plurality of end-to-end channels through said ingress aggregation, wherein said inner channels comprise time-space channels provided by the logical switches of at least one inner time-shared spatial switching stage of the ingress aggregation element.

In one embodiment, said process is further iteratively performed within each egress aggregation to determine a plurality of end-to-end channels through each said egress aggregation element, wherein said inner channels comprise time-space channels provided by the logical switches of at least one inner time-shared spatial switching stage of the egress aggregation element.

In one embodiment, the step of assigning inner time-space channels available through at least one inner time-shared spatial switching stage of the switch arrangement to said plurality of aggregated time-space channels is implemented using N processors, where N is the number of aggregations of ingress or egress elements of the switch arrangement.

In one embodiment, each ingress subelement of the switch arrangement is associated with a plurality of processors arranged to operate in parallel with each other, each processor finding a number of available channels between an ingress subelement and a switch in the final switching stage of an ingress aggregation, wherein the channels within the ingress aggregation element are found by sequential inspection of the status of channels within the ingress aggregation element.

In one embodiment, each egress subelement of the switch arrangement is associated with a plurality of processors arranged to operate in parallel with each other, each processor finding a number of available channels between a switch in the first switching stage of an egress aggregation and an egress subelement of the switch arrangement, wherein the channels within the egress aggregation element are found by sequential inspection of the status of channels within the egress aggregation element.

In one embodiment, each ingress aggregation of the switch arrangement is provided with one or more processors arranged to operate in parallel with the processors of the other ingress aggregations of the switch arrangement, each processor finding the required number of available aggregated channels between an ingress aggregation and an egress aggregation via the a switch in the final switching stage of an ingress aggregation, wherein the channels are found by sequential inspection.

In one embodiment, each ingress aggregation of the switch arrangement is associated with a plurality of processors arranged to operate in parallel with each other, each processor finding a number of available aggregated channels between an ingress aggregation and an egress aggregation via the inner time-shared spatial switching stage of the switch arrangement, wherein the channels are found by sequential inspection of the status of channels from the ingress aggregation and to the egress aggregation.

In one embodiment, each ingress subelement is associated with n separate processors to perform sequential searching for available channels, each processor searching sequentially through their channels in parallel with each other; the channel assignment process further comprising the steps of:

each processor counting the number of free channels it determines are available;

a processor sequentially adding the counts from each individual processor until the sum reaches or exceeds the required number of channels or the summation is completed.

In one embodiment, each ingress aggregation is associated with n separate processors to perform sequential searching for available aggregated channels, each processor searching sequentially through their aggregated channels in parallel with each other; the channel assignment process further comprising the steps of:

each processor counting the number of free aggregated channels it determines are available;

a processor sequentially adding the counts from each individual processor until the sum reaches or exceeds the required number of aggregated channels or the summation is completed.

In one embodiment, each ingress subelement is associated with n separate processors to perform sequential searching for available channels, each processor searching sequentially through their channels in parallel with each other; the channel assignment process further comprising the steps of:

each processor counting the number of free channels it determines are available; and
a processor or processors adding the summations from the n processors in parallel.

In one embodiment, each ingress aggregation is associated with n separate processors to perform sequential searching for available aggregated channels, each processor searching sequentially through their aggregated channels in parallel with each other; the channel assignment process further comprising the steps of:

each processor counting the number of free aggregated channels it determines are available; and

a processor or processors adding the summations from the n processors in parallel.

In one embodiment, the summations from the n processors in parallel are added using multiple stages of a real or a virtual tree of processors.

In one embodiment, the summations from the n processors in parallel are added using $\log_2 n$ stages of a real or a virtual binary tree of processors.

In one embodiment, a plurality of egress subelements receiving traffic with the same destination are treated as a single logical entity.

According to another aspect of the invention there is provided a scheduling process for a switch arrangement, the switch arrangement being arranged to perform the switching of traffic in both the space and time domains, the process comprising: receiving traffic along a number of ingress subelements of the switch; matching traffic at an ingress subelement of the switch to one or more available egress subelements of the switch; and assigning channels to matched traffic by performing the steps of: receiving traffic along a number of ingress subelements of the switch; aggregating one or more time-slots from each of said plurality of ingress subelements to form a plurality of time-space channels which are aggregated by an aggregation of the said plurality of ingress subelements; assigning inner time-space channels available through at least one inner time-shared spatial switching stage of the switch arrangement to said plurality of aggregated time-space channels; and assigning a plurality of end-to-end time-space channels from the plurality of ingress subelements to a plurality of egress subelements through the switch

arrangement after said aggregated channels have been assigned for switching through the said at least one inner time-shared spatial switching stage.

In one embodiment, the number of time-space inner channels provided by the logical switches of each of said at least one inner time-shared spatial switching stage of the switch arrangement to each aggregation element is equal to at least the number of time-slots that each ingress subelement switches end to end through the switch arrangement multiplied by the number of ingress subelements which form an aggregation element.

Another aspect of the invention comprises a switching process comprising performing steps in a method aspect, and further comprising the step of: switching received traffic along said assigned end-to-end channels from the ingress subelements to the egress subelements.

In one embodiment, traffic received by an ingress subelement is switched at least once in the time-domain within an ingress aggregation comprising said ingress subelement before being switched spatially by at least one inner time-shared spatial switching stage of the switch arrangement.

According to another aspect of the invention, there is provided a switch arrangement arranged to perform the switching of traffic in both the space and time domains, the switch arrangement comprising: a plurality of ingress subelements; a plurality of egress subelements; means to receive traffic along said number of ingress subelements; means to store said received traffic; means to aggregate one or more time-slots from each of said plurality of ingress subelements to

form a plurality of time-space channels which are aggregated for an ingress aggregation of the said plurality of ingress subelements; and at least one inner time-shared spatial switching stage, comprising a plurality of switches, wherein said plurality of logical switches are arranged to provide a number of time-space inner channels to each aggregation element which is equal to at least the number of time-slots that each ingress subelement switches end to end through the switch arrangement multiplied by the number of ingress subelements which form an aggregation element; and switch processor means arranged to switch traffic received by the switch arrangement along a plurality of assigned end-to-end channels from the ingress subelements to the egress subelements.

In one embodiment, traffic received by an ingress subelement is switched at least once in the time-domain within an ingress aggregation comprising said ingress subelement before being switched spatially by at least one inner time-shared spatial switching stage of the switch arrangement.

In one embodiment, the switch arrangement further comprises: means to assign inner time-space channels available through at least one inner time-shared spatial switching stage of the switch arrangement to said plurality of aggregated time-space channels; and means to assign a plurality of end-to-end time-space channels from the plurality of ingress subelements to a plurality of egress subelements through the switch arrangement after said aggregated channels have been assigned for switching through the said at least one inner time-shared spatial switching stage.

In one embodiment, the switch is arranged to perform frame-based switching.

In one embodiment, the switch arrangement is provided with storage means to queue traffic at its subelements.

In one embodiment, the storage means is implemented using virtual output queuing.

In one embodiment, the virtual output queues are implemented in random access memory.

In an embodiment, the switch arrangement is arranged to switch cells or packets.

In an embodiment, the switch arrangement includes at least one wavelength switch.

In an embodiment, the ingress subelements and egress subelements are bi-directional.

In an embodiment, the number of ingress subelements is not equal to the number of egress subelements.

In an embodiment, the switch is asymmetric in that the number of ingress aggregations is not equal to the number of egress aggregations.

In an embodiment, the switch arrangement comprises a multi-stage switching structure, wherein within each aggregation of ingress subelements, at least one switching stage is provided.

In an embodiment, the switch arrangement comprises a multi-stage switching structure, wherein within each aggregation of egress subelements, at least one switching stage is provided.

In an embodiment, the switch arrangement comprises a multi-stage switching structure, wherein within each aggregation of ingress subelements, at least one time-switching stage and/or at least one spatial switching stage is provided.

In an embodiment, the switch arrangement comprises a multi-stage switching structure, wherein within each aggregation of egress subelements, at least one time-switching stage and/or at least one spatial switching stage is provided.

In an embodiment, the switch arrangement comprises a multi-stage switching structure, wherein within at least one ingress aggregation of ingress subelements, a first time-switching stage; a spatial switching stage; and a second time-switching stage are provided.

In an embodiment, the switch arrangement comprises a multi-stage switching structure, wherein within at least one aggregation of egress subelements, a first time-switching stage; a spatial switching stage; and a second time-switching stage are provided.

In an embodiment, at least one inner spatial switching stage of the switch arrangement comprises one or more space switches which are shared in time between the ingress and egress aggregation elements.

In an embodiment, traffic is logically switched by the switch arrangement firstly in a time-switching stage within an aggregation element, secondly by a spatial switching stage within an aggregation element, thirdly by a time-switching stage within an aggregation element, fourthly by a time-shared spatial switching stage of the switch arrangement, fifthly in a time-switching stage within an aggregation element, sixthly by a spatial switching stage within an aggregation element, and finally by a time-switching stage within an aggregation element.

In an embodiment, the switch arrangement comprises a plurality of switching stages forming a permutation of said seven time and spatial switching stages.

In an embodiment, the arrangement of one or more time-switching stages provided within each ingress aggregation of ingress subelements and/or each egress aggregation of egress subelements implements one or more of the following:

the prevention of data contention at both the ingress subelements and the egress subelements of the switch arrangement;

the correction of sequencing of data at the egress subelements of the switch arrangement; and

the provision of contiguity of data at the egress subelements of the switch arrangement.

In an embodiment, at least one time-switching stage within is implemented using one or more time-slot interchangers.

In an embodiment, at least one time-switching stage within an ingress aggregation is implemented using one or more virtual output queues implemented in random access memory.

In an embodiment, the switch arrangement includes parallel processor means arranged to assign said plurality of aggregated time-space channels in parallel to said logical inner channels through the switch arrangement.

In an embodiment, the switch arrangement includes parallel processor means arranged to assign a plurality of time-space channels in parallel to said logical inner channels through an ingress and/or egress aggregation.

In an embodiment, the parallel processor means comprises a tree of processors.

In an embodiment, the switch arrangement is arranged to implement a channel assignment process according to a channel assignment aspect of the invention.

Other aspects of the invention include a network comprising one or more switch arrangements according to the switch arrangement aspect, a suite of one or more computer programs arranged to implement the channel assignment process aspect, a suite of at least one computer programs arranged to implement the scheduling process aspect, a suite of one or more computer programs arranged to implement the switching process aspect.

In an embodiment, the suite of at least one computer program aspect is at least partly arranged to be implemented in hardware.

Another aspect of the invention provides a scheduling process aspect including the channel assignment process aspect.

Another aspect of the invention provides a switching process aspect in which traffic is scheduled through a switch arrangement aspect using a scheduling process aspect.

Another aspect of the invention comprises a channel assignment process for a multi-stage switch arrangement having a plurality of inputs arranged in a plurality of logical associations and a plurality of outputs, wherein time-slotted traffic is received by each logical association of inputs is operated on by one or more switching stages arranged to operate only on traffic provided by the respective logical association of inputs, the channel assignment process comprising:

for each logical association, aggregating the time-slots carrying traffic from the inputs forming said logical association to form a channel comprising a plurality of logically associated time-slots;

determining a path through a spatial switching stage of the switch arrangement arranged to receive a said channel from each of said logical associations of inputs of the switch arrangement; and

determining a path for each time-slot within each logical association such that a plurality of end-to-end time-space channels are provided for one or more inputs of the switch arrangement

to their requested output via said channel through said spatial switching stage of the switch arrangement.

In an embodiment, the outputs of the switch arrangement are logically associated with one or more switching stages, and said step of determining a path for each time-slot within each logical association further comprising determining a path within each logical association of the outputs of the switch arrangement such that said plurality of end-to-end time-space channels are provided.

As set out in the appended claims, the channel assignment process provides a process for assigning channels for the service requests of a switch arrangement, and more specifically a process for assigning time-slots for the service requests of a switch arrangement. For example, a service request may comprise a request for bandwidth if the switch is a circuit-switch arrangement or a request for bit rate. In other examples, where a packet or cell switch arrangement is provided, however, a service request is typically a request for one or more time-slots to carry a predetermined number of packets or cells from an input to a predetermined output of the switch arrangement.

Another aspect of the invention seeks to provide a multi-stage time-slot assignment process for an input-queued switch arrangement in a communications network, the switch comprising a plurality of N ingress elements and N egress elements, each of the ingress elements having a number L of ingress subelements and each of the egress elements having a plurality L of egress subelements, the switch arrangement being arranged to have L or more real middle stage space

switches of size $N \times N$, and having F or more time-slots, the time-slot assignment process between the said ingress subelements and egress subelements comprising the steps of: aggregating F or more time slots from each of a plurality L in number of said ingress subelements to form an ingress element having a plurality LF or more in number of time-space channels which are pooled between the L subelements of each ingress element and the L subelements of each egress element, wherein each time-space channel corresponds to a different logical middle-stage switch of the packet switch arrangement so that the number of logical input elements and logical output elements for which channel assignment is performed through the middle stage of the switch is N ; performing time-space channel assignment through the middle stage of the switch between the logical input elements and the logical output elements; providing time-slot interchange capabilities at the 1st, 3rd, 5th stages; and performing time-slot assignment between the ingress subelements of the ingress elements and the logical output ports of the 3rd stage switches through the 2nd stage time-shared space switches, and performing time-slot assignment between the logical input ports of the 5th stage switches and the egress subelements of the egress elements through the 6th stage time-shared space switches.

Preferably, the switch arrangement is arranged to switch cells or packets. Alternatively, the switch arrangement is a circuit switch. Preferably, the switch arrangement is a cell switch arrangement capable of switching packets. Preferably, the switch arrangement further comprises a 7th stage providing a time-slot interchange capability. Preferably, the subelements comprise ports, and the elements comprise aggregations of ports. Preferably, the input-queued switch comprises VOQs which are implemented in random access memory RAM. Preferably, the time-

slot assignment process further comprises the recursive decomposition of the three stages in each element into seven stages using the steps indicated in the above aspect.

Another aspect of the invention seeks to provide a channel assignment process for a multi-stage switch arrangement having a plurality of inputs arranged in a plurality of logical associations and a plurality of outputs, wherein time-slotted traffic is received by each logical association of inputs is operated on by one or more switching stages arranged to operate only on traffic provided by the respective logical association of inputs, the channel assignment process comprising: for each logical association, aggregating the time-slots carrying traffic from the inputs forming said logical association to form a channel comprising a plurality of logically associated time-slots; determining a path through a spatial switching stage of the switch arrangement arranged to receive a said channel from each of said logical associations of inputs of time slots from each of a plurality L in number of said ingress subelements to form an ingress element having a plurality LF or more in number of time-space channels which are pooled between the L subelements of each ingress element and the L subelements of each egress element, wherein each time-space channel corresponds to a different logical middle-stage switch of the packet switch arrangement so that the number of logical input elements and logical output elements for which channel assignment is performed through the middle stage of the switch is N; performing time-space channel assignment through the middle stage of the switch between the logical input elements and the logical output elements; providing time-slot interchange capabilities at the 1st, 3rd, 5th stages; and performing time-slot assignment between the ingress subelements of the ingress elements and the logical output ports of the 3rd stage switches through the 2nd stage time-shared space switches, and performing time-slot assignment between the

HILL et al
Appl. No. 10/573,250
U.S. National Phase of PCT/GB2004/004148
August 2, 2006

logical input ports of the 5th stage switches and the egress subelements of the egress elements through the 6th stage time-shared space switches.